

AMENDMENTS TO THE CLAIMS

The following listing of claims will replace all prior versions, and listings, of claims in the application.

- 1 1-20 (Cancelled).
- 1 21. (New) A method for performing a frequent itemset operation, the method
2 comprising the steps of:
3 performing the frequent itemset operation in a plurality of phases, wherein each phase
4 is associated with combinations that have a particular number of items;
5 during at least one phase of the plurality of phases, performing the steps of
6 determining candidate combinations that are to be evaluated during the phase;
7 grouping the candidate combinations into clusters, wherein each cluster
8 corresponds to a common combination of items, and wherein all
9 candidate combinations in a given cluster include the common
10 combination of items associated with the cluster; and
11 processing said candidate combinations, based on said clusters, to determine
12 whether the candidate combinations satisfy a frequency criteria
13 associated with said frequent itemset operation.
- 1 22. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 21.
- 1 23. (New) The method of Claim 21, wherein the step of grouping the candidate
2 combinations into clusters includes the step of establishing an ordering for said
3 candidate combinations by sorting the candidate combinations relative to each other
4 based on the items within each of the candidate combinations.
- 1 24. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 23.

1 25. (New) The method of Claim 23, wherein the step of processing the candidate
2 combinations based on the clusters includes processing the candidate combinations in
3 a sequence based on said ordering.

1 26. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 25.

1 27. (New) The method of Claim 21, wherein the step of grouping the candidate
2 combinations into clusters includes hashing the candidate combinations into buckets
3 based on the items that the candidate combination contain.

1 28. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 27.

1 29. (New) The method of Claim 21, wherein the step of processing the candidate
2 combinations includes generating bitmaps for the candidate combinations, and
3 determining how many item groups of an item group population include each
4 candidate combination based on the bitmap for the candidate combination.

1 30. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 29.

1 31. (New) The method of Claim 29, wherein the step of processing the candidate
2 combinations includes, for each cluster, performing the steps of:
3 generating a bitmap for a particular combination that is a subcombination of all
4 combinations in the cluster;
5 using the bitmap for the particular combination to generate bitmaps for all
6 combinations in the cluster;
7 using the bitmap generated for each combination in the cluster to determine how
8 many item groups include the combination; and

9 after all combinations in the cluster have been processed, discarding from volatile
10 memory the bitmap for the particular combination.

1 32. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 31.

1 33. (New) The method of Claim 21, wherein the step of processing the candidate
2 combinations includes generating bitmaps for the candidate combinations as the
3 candidate combinations are processed in a sequence, the method further comprising
4 the steps of:

5 generating one or more intermediary bitmaps for use in generating of a bitmap for a
6 current candidate combination; and

7 after generating the bitmap for the current candidate combination, retaining in volatile
8 memory only those intermediary bitmaps that are base bitmaps of a next
9 candidate combination in said sequence; and

10 if any intermediate bitmaps are retained, then using one or more of the intermediary
11 bitmaps to generate a bitmap for the next candidate combination in said
12 sequence.

1 34. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 33.

1 35. (New) A method for performing a frequent itemset operation, the method
2 comprising the steps of:
3 performing the frequent itemset operation in a plurality of phases, wherein each phase
4 is associated with combinations that have a particular number of items;
5 during at least one phase of the plurality of phases, performing the steps of
6 determining candidate combinations that are to be evaluated during the phase;
7 processing said candidate combinations to determine whether the candidate
8 combinations satisfy a frequency criteria associated with said frequent
9 itemset operation, wherein the step of processing the candidate

combinations includes generating bitmaps for the candidate combinations; and
using an index on non-volatile memory to store a set of bitmaps that are generated during said at least one phase; and
during a subsequent phase of said plurality of phases, performing the steps of retrieving bitmaps from said index into volatile memory; and
using the bitmaps retrieved from said index to generate bitmaps for candidate combinations of said subsequent phase.

- 1 36. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 35.
- 1 37. (New) The method of Claim 35, wherein the step of using an index on non-volatile
2 memory to store a set of bitmaps includes using an index that uses the combination
3 associated with a bitmap as an index key for determining where within the index to
4 place an entry for the bitmap.
- 1 38. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 37.
- 1 39. (New) The method of Claim 35, wherein:
2 the at least one phase is a phase associated with N-item combinations; and
3 the set of bitmaps includes bitmaps associated with all N-item combinations that
4 satisfy the frequency criteria.
- 1 40. (New) A computer-readable medium carrying one or more sequences of instructions
2 which, when executed by one or more processors, causes the one or more processors
3 to perform the method recited in Claim 39.